Capacities and Capacity-Achieving Decoders for Various Fingerprinting Games

Thijs Laarhoven*

April 3, 2014

arXiv:1401.5688v3 [cs.IT] 2 Apr 2014

Abstract

Combining an information-theoretic approach to fingerprinting with a more constructive, statistical approach, we derive new results on the fingerprinting capacities for various informed settings, as well as new log-likelihood decoders with provable code lengths that asymptotically match these capacities. The simple decoder built against the interleaving attack is further shown to achieve the simple capacity for unknown attacks, and is argued to be an improved version of the recently proposed decoder of Oosterwijk et al. With this new universal decoder, cut-offs on the bias distribution function can finally be dismissed.

Besides the application of these results to fingerprinting, a direct consequence of our results to group testing is that (i) a simple decoder asymptotically requires a factor 1.44 more tests to find defectives than a joint decoder, and (ii) the simple decoder presented in this paper provably achieves this bound.

1 Introduction

To protect copyrighted content against unauthorized redistribution, distributors may embed watermarks or fingerprints in the content, uniquely linking copies to individual users. Then, if an illegal copy of the content is found, the distributor can extract the watermark from the copy and compare it to the database of watermarks, to determine which user was responsible.

To combat this solution, pirates may try to form a coalition of several colluders, each owning a differently watermarked copy of the content, and perform a collusion attack. By comparing their different versions of the content, they will detect differences in their copies which must be part of the watermark. They can then create a new pirate copy, where the resulting watermark matches the watermark of different pirates in different segments of the content, making it hard for the distributor to find the responsible users. Fortunately, under the assumption that if the pirates don't detect any differences (because they all received the same version) they output this watermark (known in the literature as the Boneh-Shaw marking assumption [5]), it is still possible to find all colluders using suitable fingerprinting codes.

*T. Laarhoven is with the Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands.
E-mail: mail@thijs.com.

1.1 Model

The above fingerprinting game is often modeled as the following two-person game between the distributor $\mathcal{D}$ and the coalition of pirates $\mathcal{C}$. The set of colluders is assumed to be a random subset of size $|\mathcal{C}| = c$ from the complete set of n users $\mathcal{U}$, and the identities of these colluders are unknown to the distributor. The aim of the game for the distributor is ultimately to discover the identities of the colluders, while the colluders want to stay hidden. The game consists of the following three phases: (i) the distributor uses an *encoder* to generate the fingerprints; (ii) the colluders employ a *collusion channel* to generate the pirate output, and (iii) the distributor uses a *decoder* to map the pirate output to a set of accused users.

Encoder First, the distributor generates a fingerprinting code $\mathcal{X}$, consisting of n code words $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_n$ from $\{0,1\}^\ell$.[1] The ith entry of code word j indicates which version of the content is assigned to user j in the ith segment. The parameter ℓ is referred to as the code length, and the distributor would like ℓ to be as small as possible.

A common restriction on the encoding process is to assume that $\mathcal{X}$ is created by first generating a probability vector $\boldsymbol{P} \in [0,1]^\ell$ by choosing each entry P_i independently from a certain distribution function F, and then generating $\mathcal{X}$ according to $\mathbb{P}(X_{j,i} = 1) = P_i$. This guarantees that watermarks of different users j are independent, and that watermarks in different positions i are independent. Schemes that satisfy this assumption are sometimes called *bias-based schemes*, and the encoders discussed in this paper also belong to this category.

Collusion channel After generating $\mathcal{X}$, the entries are used to select and embed watermarks in the content, and the content is sent out to all users. The colluders then get together, compare their copies, and use a certain collusion channel or pirate attack $\boldsymbol{\theta}$ to select the pirate output $\boldsymbol{Y} \in \{0,1\}^\ell$. If the pirate attack behaves symmetrically both in the colluders and in the positions i, then the collusion channel can be modeled by a vector $\boldsymbol{\theta} \in [0,1]^{c+1}$, consisting of entries $\theta_z = f_{Y|Z}(1|z)$ indicating the probability of outputting a 1 when the pirates received z ones and $c - z$ zeroes. Some common attacks $\boldsymbol{\theta}$ are described in Section 2.3.

[1] More generally $\mathcal{X}$ is a code with code words of length ℓ from an alphabet $\mathcal{Q}$ of size $q \geq 2$, but in this paper we restrict our attention to the binary case $q = 2$.

Decoder Finally, after the pirate output has been generated and distributed, we assume that the distributor intercepts it and applies a decoding algorithm to the pirate output $\boldsymbol{Y}$, the code $\mathcal{X}$ and the (secret) bias vector $\boldsymbol{P}$ to compute a set $\mathcal{C}' \subseteq \mathcal{U}$ of accused users. This is commonly done by assigning *scores* to users, and accusing those users whose score exceeds some predefined threshold η. The distributor wins the game if $\mathcal{C}'$ is non-empty and contains only colluders (i.e. $\emptyset \neq \mathcal{C}' \subseteq \mathcal{C}$) and loses if this is not the case, which could be because an innocent user $j \notin \mathcal{C}$ is falsely accused (a false positive error), or because no guilty users are accused (a false negative error). We often write ε_1 and ε_2 for upper bounds on the false positive and false negative probabilities respectively.

1.2 Related work

Work on the above bias-based fingerprinting game started in 2003, when Tardos proved that any fingerprinting scheme must satisfy $\ell \propto c^2 \ln n$, and that a bias-based scheme is able to achieve this optimal scaling in ℓ [38]. He proved the latter by providing a simple and explicit construction with a code length of $\ell = 100c^2 \ln(n/\varepsilon_1)$, which is known in the literature as the Tardos scheme.

Improved constructions Later work on the constructive side of fingerprinting focused on improving upon Tardos' result by sharpening the bounds [3, 35], optimizing the distribution functions [27], improving the score function [36], tightening the bounds again with this improved score function [18, 22, 28, 34, 36, 37], optimizing the score function [29], and again tightening the bounds with this optimized score function [16, 30] to finally end up with a sufficient asymptotic code length of $\ell \sim 2c^2 \ln n$ for large n. This construction can be extended to larger alphabets, in which case the code length scales as $\ell \sim 2c^2 \ln(n)/(q-1)$. Other work on practical constructions focused on joint decoders, which are computationally more involved but may work with shorter codes [24, 25, 31], and side-informed fingerprinting games [7, 10, 21, 29], where estimating the collusion channel $\boldsymbol{\theta}$ was considered to get an improved performance.

Recently Abbe and Zheng [1] showed that, in the context of fingerprinting [24], if the set of allowed collusion channels satisfies a certain one-sidedness condition, then a decoder that achieves capacity against the information-theoretic worst-case attack is a universal decoder achieving capacity against arbitrary attacks. The main drawback of using this result is that the worst-case attack is hard to compute, but this does lead to more insight why e.g. Oosterwijk et al. [30] obtained a universal decoder by considering the decoder against the 'interleaving attack', which is known to be the asymptotic worst-case attack.

Fingerprinting capacities At the same time, work was also done on establishing bounds on the fingerprinting capacity C, which translate to lower bounds on the required asymptotic code length ℓ through $\ell \gtrsim C^{-1} \log_2 n$ for large n. For the binary case Huang and Moulin [11–14, 25] and Amiri and Tardos [2] independently derived exact asymptotics for the fingerprinting capacity for arbitrary attacks as $C \sim (2c^2 \ln 2)^{-1}$, corresponding to a minimum code length of $\ell \sim 2c^2 \ln n$. Huang and Moulin [14] further showed that to achieve this bound, an encoder should use the arcsine distribution F^* for generating biases p:

$$F^*(p) = \frac{2}{\pi} \arcsin \sqrt{p}. \qquad (0 < p < 1) \qquad (1)$$

These capacity-results were later generalized to the q-ary setting [4, 15] showing that a q-ary code length of $\ell \sim 2c^2 \ln(n)/(q-1)$ is asymptotically optimal.

Dynamic fingerprinting There has also been some interest in a variant of the above fingerprinting game where several rounds of the two-player game between the distributor and the coalition are played sequentially. This allows the distributor to adjust the encoding and decoding steps of the next rounds to the knowledge obtained from previous rounds. Many of the bias-based constructions can also be used effectively in this dynamic setting [17, 20, 21] with equivalent asymptotics for the required code length, but allowing the distributor to trace all colluders even if the collusion channel is not symmetric in the colluders, and leading to significantly smaller first order terms than in the 'static' setting. These bias-based dynamic schemes may even be able to compete with the celebrated scheme of Fiat and Tassa [9].

Group testing Finally, a different area of research closely related to fingerprinting is that of group testing, where the set of n users corresponds to a set of n items, the set of c colluders corresponds to a subset of c defective items, and where the aim of the distributor is to find all defective items by performing group tests. This game corresponds to a special case of the fingerprinting game, where the pirate attack is fixed in advance (and possibly known to the distributor) to (a variant of) the 'all-1 attack'. In this game it is significantly easier to find all pirates/defectives; it is known that a joint decoder asymptotically requires only $\ell \sim c \log_2 n$ tests [33], while simple decoders exist requiring as few as $\ell \sim ec \ln n$ tests to find all defectives [6]. Recent work has shown that applying results from fingerprinting to group testing may lead to improved results compared to what is known in the group testing literature [19, 23].

1.3 Contributions

In this work we first extend the work of Huang and Moulin [14] by deriving explicit asymptotics for the simple and joint capacities of various fingerprinting games with different amounts of side-information. Table 1 summarizes tight lower bounds on the code length constant for various informed settings obtained via the capacities. These asymptotics can be seen as our 'targets' for the second part of this paper, which describes decoders with provable bounds on ℓ and η that asymptotically achieve these capacities. In fact, if the collusion channel that the decoder was built against matches the attack used by the pirates, then the proof that the resulting simple decoders achieve

Table 1: Asymptotics for tight lower bounds on $L = \ell/\ln n \cong C^{-1}/\ln 2$, based on the simple and joint capacities, with different amounts of side information (see Section 2.2). The proposed simple decoders are shown to match these bounds, and we conjecture that the proposed joint decoders are also asymptotically optimal.

	Fully informed		Partially informed	
	Simple	Joint	Simple	Joint
Interleaving atk.	$2c^2$	$2c^2$	$2c^2$	$2c^2$
All-1 attack	$2.08c$	$1.44c$	$1.83c\sqrt{c}$	$1.32c\sqrt{c}$
Majority voting	$3.14c$	$1.44c$	$2.41c\sqrt{c}$	$1.20c\sqrt{c}$
Minority voting	$2.08c$	$1.44c$	$0.66c\sqrt{c}$	$0.43c\sqrt{c}$
Coin-flip attack	$8.33c$	$4.48c$	$5.18c\sqrt{c}$	$2.32c\sqrt{c}$

capacity is remarkably simple and holds for arbitrary attacks.

Capacity-achieving simple decoding without cut-offs Similar to Oosterwijk et al. [29,30], who studied the decoder built against the interleaving attack because that attack is in a sense optimal, we then turn our attention to the simple decoder designed against the interleaving attack, and argue that it is an improved version of Oosterwijk et al.'s universal decoder. To provide a sneak preview of this result, the new score function is the following:

$$g(x,y,p) = \begin{cases} \ln\left(1+\frac{p}{c(1-p)}\right) & x=y=0 \\ \ln\left(1-\frac{1}{c}\right) & x \neq y \\ \ln\left(1+\frac{1-p}{cp}\right) & x=y=1 \end{cases} \tag{2}$$

This decoder is shown to achieve the uninformed simple capacity, and we argue that with this decoder (i) the Gaussian assumption always holds (and convergence to the normal distribution is much faster), and (ii) no cut-offs on the bias distribution function F are ever needed anymore.

Joint log-likelihood decoders Since it is not hard to extend the definition of the simple decoder to joint decoding, we also present and analyze joint log-likelihood decoders. Analyzing these joint decoders turns out to be somewhat harder due to the 'mixed tuples', but we give some motivation why these decoders seem to work well. We also conjecture that the joint decoder tailored against the interleaving attack achieves the joint uninformed capacity, but proving this result is left for future work.

Applications to group testing Since the all-1 attack in fingerprinting is equivalent to a problem known in the literature as group testing [21, 23], some of our results can also be applied to this area. In fact, we derive two new results in the area of group testing: (i) any simple-decoder group testing algorithm requires at least $\ell \sim \log_2(e)^2 c \ln n \approx 2.08 c \ln n$ group tests to find c defective items hidden among n items, and (ii) the decoder discussed in Section 4.1 provably achieves this optimal scaling in ℓ. This decoder was previously considered in [23], but no provable bounds on the (asymptotic) code lengths were given there.

1.4 Outline

The outline of the paper is as follows. Section 2 first describes the various different models we consider in this paper, and provides a roadmap for Sections 3 and 4. Section 3 discusses capacity results for each of these models, while Section 4 discusses decoders which aim to match the lower bounds on ℓ obtained in Section 3. Finally, in Section 5 we conclude with a brief discussion of the most important results and remaining open problems.

2 Different models

Let us first describe how the results in Sections 3 and 4 are structured according to different assumptions, leading to different models. Besides the general assumptions on the model discussed in the introduction, we further make a distinction between models based on (1) the computational complexity of the decoder, (2) the information about $\boldsymbol{\theta}$ known to the distributor, and (3) the collusion channel used by the pirates. These are discussed in Sections 2.1, 2.2 and 2.3 respectively.

2.1 Decoding complexity

Commonly two types of decoders are considered, which use different amounts of information to decide whether a user should be accused or not.

1. **Simple decoding**: To quote Moulin [25, Section 4.3]: *"The receiver makes an innocent/guilty decision on each user independently of the other users, and there lies the simplicity but also the suboptimality of this decoder."* In other words, the decision to accuse user j depends only on the jth code word of $\mathcal{X}$, and not on other code words from $\mathcal{X}$.

2. **Joint decoding**: In this case, the decoder is allowed to base the decision whether to accuse a user on the entire code $\mathcal{X}$. Such decoders may be able to obtain smaller code lengths than possible with the best simple decoders.

Using more information generally causes the time complexity of the decoding step to go up, so usually there is a trade-off between a shorter code length and a faster decoding algorithm.

2.2 Side-informed distributors

We consider three different scenarios with respect to the knowledge of the distributor about the collusion channel $\boldsymbol{\theta}$. Depending on the application, different scenarios may apply.

1. **Fully informed**: Even before $\mathcal{X}$ is generated, the distributor already knows exactly what the pirate attack $\boldsymbol{\theta}$ will be. This information can thus be used to optimize both the encoding and decoding phases. This scenario applies to various group testing models, and may apply to dynamic traitor tracing, where after several rounds the distributor may have estimated the pirate strategy.

2. **Partially informed**: The tracer does not know in advance what collusion channel will be used, so the encoding is aimed at arbitrary attacks. However, after obtaining the pirate output $\boldsymbol{y}$, the distributor does learn more about $\boldsymbol{\theta}$ before running an accusation algorithm, e.g. by estimating the attack based on the available data. So the encoding is uninformed, but we assume that the decoder is informed and knows $\boldsymbol{\theta}$. Since the asymptotically optimal bias distribution function F in fingerprinting is known to be the arcsine distribution F^*, we will assume that F^* is used for generating biases. This scenario is similar to EM decoding [7, 10].

3. **Uninformed**: In this case, both the encoding and decoding phases are assumed to be done without prior knowledge about $\boldsymbol{\theta}$, so also the decoder should be designed to work against arbitrary attacks. This is the most commonly studied fingerprinting game.

For simplicity of the analysis, in the partially informed setting we assume that the estimation of the collusion channel is precise, so that $\boldsymbol{\theta}$ is known exactly to the decoder. This assumption may not be realistic, but at least we can then obtain explicit expressions for the capacities, and get an idea of how much estimating the strategy may help in reducing the code length. This also allows us to derive explicit lower bounds on ℓ: even if somehow the attack can be estimated correctly, then the corresponding capacities tell us that we will still need at least a certain number of symbols to find the pirates.

2.3 Common collusion channels

As mentioned in the introduction, we assume that collusion channels satisfy the marking assumption, which means that $\theta_0 = 0$ and $\theta_c = 1$. For the remaining values of $z \in \{1, \dots, c-1\}$ the pirates are free to choose how often they want to output a 1 when they receive z ones. Some commonly considered attacks are listed below.

1. **Interleaving attack**: The coalition randomly selects one of its members and outputs his symbol. This corresponds to $(\boldsymbol{\theta}_{\text{int}})_z = z/c$. This attack is known to be asymptotically optimal (from the point of view of the colluders) in the uninformed max-min fingerprinting game [14].

2. **All-1 attack**: The pirates output a 1 whenever they can, i.e., whenever they have at least one 1. This translates to $(\boldsymbol{\theta}_{\text{all1}})_z = \mathbb{1}\{z > 0\}$. This attack is of particular interest due to its relation with group testing.

3. **Majority voting**: The colluders output the most common symbol among their received symbols. This means that $(\boldsymbol{\theta}_{\text{maj}})_z = \mathbb{1}\{z > c/2\}$.

4. **Minority voting**: The traitors output the symbol which they received the least often (but received at least once). For $1 \le z \le c-1$, this corresponds to $(\boldsymbol{\theta}_{\min})_z = \mathbb{1}\{z < c/2\}$.

5. **Coin-flip attack**: If the pirates receive both symbols, they flip a fair coin to decide which symbol to output. So for $1 \le z \le c-1$, this corresponds to $(\boldsymbol{\theta}_{\text{coin}})_z = \frac{1}{2}$.

For even c, defining $\theta_{c/2}$ in a consistent way for majority and minority voting is not straightforward. For simplicity, in the analysis of these two attacks we will therefore assume that c is odd. Note that in the uninformed setting, we do not distinguish between different collusion channels; the encoder and decoder should then work against arbitrary attacks.

2.4 Roadmap

The upcoming two sections about capacities (Section 3) and decoders (Section 4) are structured according to the above classification, where first the decoding complexity is chosen, then the side-information is fixed, and finally different attacks are considered. For instance, to find the joint capacity in the fully informed game one has to go to Section 3.2.1, while the new simple uninformed decoder can be found in Section 4.1.3.

3 Capacities

In this section we establish lower bounds on the code length of any valid decoder, by inspecting the information-theoretic capacities of the various fingerprinting games. We will use some common definitions from information theory, such as the binary entropy function $h(x) = -x\log_2 x - (1-x)\log_2(1-x)$, the relative entropy or Kullback-Leibler divergence $d(x\|y) = x\log_2(x/y) + (1-x)\log_2((1-x)/(1-y))$, and the mutual information $I(X;Y) = \sum_{x,y} \mathbb{P}(x,y)\log_2(\mathbb{P}(x,y)/\mathbb{P}(x)\mathbb{P}(y))$. The results in this section build further upon previous work on this topic by Huang and Moulin [14].

3.1 Simple capacities

For simple decoders, we assume that the decision whether to accuse user j is based solely on $\boldsymbol{X}_j$, $\boldsymbol{Y}$ and $\boldsymbol{P}$. Focusing on a single position, and denoting the random variables corresponding to a colluder's symbol, the pirate output, and the bias in this position by X_1, Y and P, the interesting quantity to look at [11] is the mutual information $I(X_1; Y|P = p)$. This quantity depends on the pirate strategy $\boldsymbol{\theta}$ and on the bias p. To study this mutual information we will use the following equality [14, Equation (61)],

$$I(X_1; Y|P = p) = pd(a_1\|a) + (1-p)d(a_0\|a), \qquad (3)$$

where a, a_0, a_1 are defined as

$$a = \sum_{z=0}^{c} \binom{c}{z} p^z (1-p)^{c-z} \theta_z, \tag{4}$$

$$a_0 = \sum_{z=0}^{c-1} \binom{c-1}{z} p^z (1-p)^{c-z-1} \theta_z, \tag{5}$$

$$a_1 = \sum_{z=1}^{c} \binom{c-1}{z-1} p^{z-1} (1-p)^{c-z} \theta_z. \tag{6}$$

Note that given p and $\boldsymbol{\theta}$, the above formulas allow us to compute the associated mutual information explicitly.

3.1.1 Fully informed

In the fully informed setting we are free to choose F to maximize the capacity, given a collusion channel $\boldsymbol{\theta}$. When the attack is known to the distributor in advance, there is no reason to use different values of p; the distributor should always use the value of p that maximizes the mutual information payoff $I(X_1; Y|P = p)$. Given an attack strategy $\boldsymbol{\theta}$, the capacity we are interested in is thus

$$C^s(\boldsymbol{\theta}) = \max_p I(X_1; Y|P = p). \tag{7}$$

For general attacks finding the optimal value of p analytically can be hard, but for certain specific attacks we can investigate the resulting expressions individually to find the optimal values of p that maximize the mutual information. This leads to the following results for the five attacks listed in Section 2.3. Proofs will appear in the full version.

Theorem 1. *The simple informed capacities and the corresponding optimal values of p for the five attacks of Section 2.3 are:*

$$C^s(\boldsymbol{\theta}_{int}) \sim \frac{1}{2c^2 \ln 2}, \qquad p^s_{int} = \frac{1}{2}, \tag{S1}$$

$$C^s(\boldsymbol{\theta}_{all1}) \sim \frac{\ln 2}{c}, \qquad p^s_{all1} \sim \frac{\ln 2}{c}, \tag{S2}$$

$$C^s(\boldsymbol{\theta}_{maj}) \sim \frac{1}{\pi c \ln 2}, \qquad p^s_{maj} = \frac{1}{2}, \tag{S3}$$

$$C^s(\boldsymbol{\theta}_{min}) \sim \frac{\ln 2}{c}, \qquad p^s_{min} \sim \frac{\ln 2}{c}, \tag{S4}$$

$$C^s(\boldsymbol{\theta}_{coin}) \sim \frac{\ln 2}{4c}, \qquad p^s_{coin} \sim \frac{\ln 2}{2c}. \tag{S5}$$

Since fully informed protection against the all-1 attack is equivalent to noiseless group testing, and since the code length ℓ scales in terms of the capacity C as $\ell \geq C^{-1} \log_2 n$, we immediately get the following corollary.

Corollary 1. *Any simple group testing algorithm for c defectives and n total items requires an asymptotic number of group tests ℓ of at least*

$$\ell \sim \frac{c \ln n}{\ln(2)^2} \approx 2.08 c \ln n. \tag{8}$$

Note that this seems to contradict earlier results of [19], which suggested that under a certain Gaussian assumption, only $\ell \sim 2c \ln n$ tests are required. This apparent contradiction is caused by the fact that the Gaussian assumption in [19] is not correct in the regime of small p, for which those results were derived. In fact, the distributions considered in that paper roughly behave like binomial distributions over ℓ trials with probability of success of $O(1/\ell)$, which converge to Poisson distributions. Numerical inspection shows that the relevant distribution tails are indeed not very Gaussian and do not decay fast enough. Rigorous analysis of the scores in [19] shows that an asymptotic code length of about $3c \ln n$ is sufficient when $p \sim \ln(2)/c$, which is well above the lower bound of Corollary 1. Details can be found in the full version.

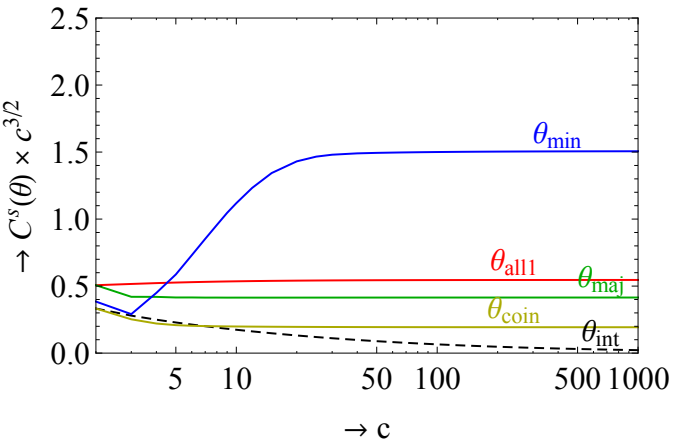


Figure 1: The partially informed simple capacity (multiplied by $c^{3/2}$) as a function of c for different pirate attacks. From top to bottom these curves correspond to minority voting, the all-1 attack, majority voting, the coin-flip attack, and the interleaving attack. Except for the interleaving attack, for which the capacity scales as c^{-2} (the dashed line), these capacities all seem to scale as $c^{-3/2}$.

3.1.2 Partially informed

If the encoder is uninformed, then the best he can do against arbitrary attacks (for large c) is to generate biases using the arcsine distribution F^*. So instead of computing the mutual information in one point $P = p$, we now average over different values of p where p follows the arcsine distribution. So the capacity we are interested in is given by

$$C^s(\boldsymbol{\theta}) = \mathbb{E}_p I(X_1; Y|P = p) = \int_0^1 \frac{I(X_1; Y|P = p)}{\pi \sqrt{p(1-p)}}\, dp. \tag{9}$$

The resulting integrals are hard to evaluate analytically, even for large c, although for some collusion channels we can use Pinsker's inequality (similar to the proof of [14, Theorem 7]) to show that $C^s(\boldsymbol{\theta}) = \Omega(c^{-3/2})$. And indeed, if we look at the numerics of $c^{3/2} C^s(\boldsymbol{\theta})$ in Figure 1, it seems that the partially informed capacity usually scales as $c^{-3/2}$. As a consequence, even if the attack can be estimated exactly, then still a code length of the order $\ell \propto c^{3/2} \ln n$ is required to get a scheme that works. Note that for the interleaving attack, the capacity scales as c^{-2}.

3.1.3 Uninformed

For the uninformed fingerprinting game, where both the encoder and decoder are built to work against arbitrary attacks, we are interested in the following max-min game:

$$C^s = \max_F \min_{\boldsymbol{\theta}(F)} \mathbb{E}_p I(X_1; Y|P=p). \tag{10}$$

Huang and Moulin [14, 15] previously solved this uninformed game for asymptotically large coalition sizes c as follows.

Proposition 1. *[15, Theorem 3] The simple uninformed capacity is given by*

$$C^s \sim \frac{1}{2c^2 \ln 2}, \tag{11}$$

and the optimizing encoder F and collusion channel $\boldsymbol{\theta}$ achieving this bound for large c are the arcsine distribution F^* and the interleaving attack $\boldsymbol{\theta}_{int}$.

Note that while for the interleaving attack the capacity is the same (up to order terms) for each of the three side-informed cases, for the four other attacks the capacity gradually increases from $O(c^{-2})$ to $O(c^{-3/2})$ to $O(c^{-1})$ when the distributor is more and more informed.

3.2 Joint capacities

If the computational complexity of the decoder is not an issue, joint decoding may be an option. In that case, the relevant quantity to examine is the mutual information between the symbols of all colluders, denoted by $X_1, \ldots, X_c$, and the pirate output Y, given P: $I(X_1, \ldots, X_c; Y|P=p)$ [14]. Note that Y only depends on $X_1, \ldots, X_c$ through $Z = \sum X_i$, so $I(X_1, \ldots, X_c; Y|P=p) = I(Z; Y|P=p)$. To compute the joint capacities, we use the following convenient explicit formula [14, Equation (59)]:

$$\frac{1}{c} I(Z; Y|P=p) = \frac{1}{c}\left[h(a) - a_h\right], \tag{12}$$

where $h(p) = -p\log_2 p - (1-p)\log_2(1-p)$ is the binary entropy function, and a_h is defined as

$$a_h = \sum_{z=0}^{c} \binom{c}{z} p^z (1-p)^{c-z} h(\theta_z). \tag{13}$$

3.2.1 Fully informed

In the fully informed setting, the capacity is again obtained by considering the mutual information and maximizing it as a function of p:

$$C^j(\boldsymbol{\theta}) = \frac{1}{c} \max_p I(Z; Y|P=p). \tag{14}$$

Computing this is very easy for the all-1 attack, the majority voting attack and the minority voting attack, since one can easily prove that the joint capacity is equal to $\frac{1}{c}$ whenever the collusion channel is deterministic, e.g. when $\theta_z \in \{0, 1\}$ for all z. Since the capacity for the interleaving attack was already known, the only non-trivial case is the coin-flip attack. A proof of the following theorem can be found in the full version.

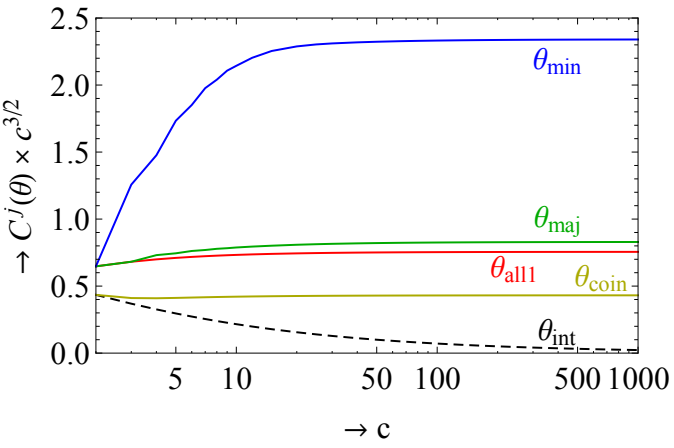


Figure 2: The partially informed joint capacity, as a function of c, for different attacks. From top to bottom these are minority voting, majority voting, the all-1 attack, the coin-flip attack and the interleaving attack. Except for the interleaving attack, these capacities all seem to scale as $c^{-3/2}$.

Theorem 2. *The joint informed capacities and the corresponding optimal values of p for the five attacks of Section 2.3 are:*

$$C^j(\boldsymbol{\theta}_{int}) \sim \frac{1}{2c^2 \ln 2}, \qquad p^j_{int} = \frac{1}{2}, \tag{J1}$$

$$C^j(\boldsymbol{\theta}_{all1}) = \frac{1}{c}, \qquad p^j_{all1} \sim \frac{\ln 2}{c}, \tag{J2}$$

$$C^j(\boldsymbol{\theta}_{maj}) = \frac{1}{c}, \qquad p^j_{maj} = \frac{1}{2}, \tag{J3}$$

$$C^j(\boldsymbol{\theta}_{min}) = \frac{1}{c}, \qquad p^j_{min} = \frac{1}{2}, \tag{J4}$$

$$C^j(\boldsymbol{\theta}_{coin}) \sim \frac{\log_2(5/4)}{c}, \qquad p^j_{coin} \sim \frac{\ln(5/3)}{c}. \tag{J5}$$

Recall that there is a one-to-one correspondence between the all-1 attack and group testing, so the result above establishes firm bounds on the asymptotic number of group tests required by any probabilistic group testing algorithm. This result was already known, and was first derived by Sebő [33, Theorem 2].

3.2.2 Partially informed

For the partially informed capacity we again average over the mutual information where p is drawn at random from the arcsine distribution F^*. Thus the capacity is given by

$$C^j(\boldsymbol{\theta}) = \frac{1}{c} \mathbb{E}_p I(Z; Y|P=p). \tag{15}$$

Exact results are again hard to obtain, but we can at least compute the capacities numerically to see how they behave. Figure 2 shows the capacities of the five attacks of Section 2.3. Although the capacities are higher for joint decoding than for simple decoding, the joint capacities of all attacks but the interleaving attack also scale as $c^{-3/2}$.

3.2.3 Uninformed

Finally, if we are working with joint decoders which are supposed to work against arbitrary attacks, then we are

interested in the following max-min mutual information game:

$$C^j = \max_F \min_{\boldsymbol{\theta}(F)} \mathbb{E}_p \frac{1}{c} I(Z; Y|P = p). \tag{16}$$

This joint capacity game was previously solved by Huang and Moulin [14] who showed that also in the joint game, the interleaving attack $\boldsymbol{\theta}_{\text{int}}$ and the arcsine distribution F^* together form a saddle-point solution to the uninformed fingerprinting game.

Proposition 2. *[14, Theorem 6, Corollary 7] The joint uninformed capacity is given by*

$$C^j \sim \frac{1}{2c^2 \ln 2}, \tag{17}$$

and the optimizing encoder F and collusion channel $\boldsymbol{\theta}$ achieving this bound for large c are the arcsine distribution F^* and the interleaving attack $\boldsymbol{\theta}_{int}$.

4 Decoders

After deriving "targets" for our decoders in the previous section, this section discusses decoders that aim to match these bounds. We will follow the score-based framework introduced by Tardos [38], which was later generalized to joint decoders by Moulin [25]. For simple decoding, this means that a user j receives a score S_j of the form

$$S_j = \sum_{i=1}^{\ell} S_{j,i} = \sum_{i=1}^{\ell} g(X_{j,i}, y_i, p_i), \tag{18}$$

where g is called the score function. User j is then accused if $S_j > \eta$ for some threshold η.

For joint decoding, scores are assigned to tuples $T = \{j_1, \ldots, j_c\}$ of c distinct users according to

$$S_T = \sum_{i=1}^{\ell} S_{T,i} = \sum_{i=1}^{\ell} g(X_{j_1,i}, \ldots, X_{j_c,i}, y_i, p_i). \tag{19}$$

In this case, a tuple of users is accused if the joint tuple score exceeds some other threshold η. Note that this accusation algorithm is not exactly well-defined, since it is possible that a user appears both in a tuple that is accused and in a tuple that is not accused. For the analysis we will assume that the scheme is only successful if the single tuple consisting of all colluders has a score exceeding η and no other tuples have a score exceeding η, in which case all users in the guilty tuple are accused.

4.1 Simple decoders

Several different score functions for the simple decoder setting were considered before, but in this work we will restrict our attention to the following log-likelihood scores, which perform well and turn out to be easy to analyze:

$$g(x, y, p) = \ln\left(\frac{\mathbb{P}_g(x, y|p)}{\mathbb{P}_i(x, y|p)}\right). \tag{20}$$

Here $\mathbb{P}_g(x, y|p) = \mathbb{P}(X_{j,i} = x, Y_i = y|P_i = p, j \in \mathcal{C})$ corresponds to the probability of seeing the pair (x, y) when user j is guilty, and $\mathbb{P}_i(x, y|p)$ corresponds to the same probability under the assumption that j is innocent. Using this score function g, the complete score of a user is the logarithm of a Neyman-Pearson score over the entire codeword:

$$S_j = \sum_{i=1}^{\ell} \ln\left(\frac{\mathbb{P}_g(x_{j,i}, y_i|p_i)}{\mathbb{P}_i(x_{j,i}, y_i|p_i)}\right) = \ln\left(\frac{\mathbb{P}_g(\boldsymbol{x}_j, \boldsymbol{y}|\boldsymbol{p})}{\mathbb{P}_i(\boldsymbol{x}_j, \boldsymbol{y}|\boldsymbol{p})}\right). \tag{21}$$

Such Neyman-Pearson scores are known to be optimally discriminative to decide whether to accuse a user or not. Log-likelihood scores were previously considered in the context of fingerprinting in e.g. [24, 32].

4.1.1 Fully informed

For the central theorem below, we will make use of the following function M, which is closely related to the moment generating functions of scores in one position i for innocent and guilty users. This function is defined as

$$M(t) = \sum_{x,y} \mathbb{P}_i(x, y|p)^{1-t} \mathbb{P}_g(x, y|p)^t \tag{22}$$

and it satisfies $M(t) = \mathbb{E}_i e^{tS_{j,i}} = \mathbb{E}_g e^{(t-1)S_{j,i}}$ and $M(0) = M(1) = 1$.

Theorem 3. *Let p and $\boldsymbol{\theta}$ be fixed and known to the distributor. Let $\gamma = \ln(1/\varepsilon_2)/\ln(n/\varepsilon_1)$, and let the threshold η and code length ℓ be defined as*

$$\eta = \ln\left(\frac{n}{\varepsilon_1}\right), \qquad \ell = \frac{\sqrt{\gamma}(1+\sqrt{\gamma})}{-\ln M(1-\sqrt{\gamma})} \ln\left(\frac{n}{\varepsilon_1}\right). \tag{23}$$

Then with probability at least $1 - \varepsilon_1$ no innocent users are accused (regardless of which collusion channel was used), and with probability at least $1 - \varepsilon_2$ a colluder is caught (if the collusion channel is indeed $\boldsymbol{\theta}$).

Proof. For innocent users j, we would like to prove that $\mathbb{P}_i(S_j > \eta) \leq \varepsilon_1/n$, where S_j is the user's total score over all positions. If this can be proved, then it follows that with probability at least $(1 - \varepsilon_1/n)^n \geq 1 - \varepsilon_1$ no innocent users are accused. Using the Markov inequality for $e^{\alpha S_j}$ with $\alpha > 0$ and optimizing over α, we see that the optimum lies close to $\alpha = 1$. For simplicity we choose $\alpha = 1$ which, combined with the given value of η, leads to the following bound:

$$\mathbb{P}_i(S_j > \eta) = \min_{\alpha>0} \mathbb{P}_i(e^{\alpha S_j} > e^{\alpha\eta}) \leq \min_{\alpha>0} \frac{\mathbb{E}_i(e^{\alpha S_j})}{e^{\alpha\eta}} \tag{24}$$

$$= \min_{\alpha>0} \frac{\prod_{i=1}^{\ell} \mathbb{E}_i(e^{\alpha S_{j,i}})}{e^{\alpha\eta}} = \min_{\alpha>0} \frac{M(\alpha)^\ell}{(n/\varepsilon_1)^\alpha} \leq \frac{M(1)^\ell}{n/\varepsilon_1} = \frac{\varepsilon_1}{n}. \tag{25}$$

For guilty users, we would like to prove that for an arbitrary guilty user j, we have $\mathbb{P}_g(S_j < \eta) \leq \varepsilon_2$. Again using

Markov's inequality (but now with a more sophisticated exponent $\beta = \sqrt{\gamma}$) we get

$$\mathbb{P}_g(S_j < \eta) \leq \min_{\beta>0} \frac{\mathbb{E}_g(e^{-\beta S_j})}{e^{-\beta\eta}} = \min_{\beta>0} \frac{\prod_{i=1}^{\ell} \mathbb{E}_g(e^{-\beta S_{j,i}})}{e^{-\beta\eta}} \quad (26)$$

$$= \min_{\beta>0} \frac{M(1-\beta)^\ell}{e^{-\beta\eta}} \leq \frac{M(1-\sqrt{\gamma})^\ell}{e^{-\sqrt{\gamma}\eta}} = \varepsilon_2, \quad (27)$$

where the last equality follows from the definitions of ℓ and η of (23). □

Compared to previous papers analyzing provable bounds on the error probabilities, the proof of Theorem 3 is remarkably short and simple. The only problem is that the given expression for ℓ is not very informative as to how ℓ scales for large n. The following corollary answers this question, by showing how ℓ scales for small γ.

Corollary 2. *If $\gamma = o(1)$ then ℓ achieves the optimal asymptotic scaling (achieves capacity) for arbitrary p:*

$$\ell = \frac{\log_2 n}{I(X_1; Y|P=p)}[1 + O(\sqrt{\gamma})], \quad (28)$$

Proof. First, let us study the behavior of $M(1-\sqrt{\gamma})$ for small γ, by computing the first order Taylor expansion of $M(1-\sqrt{\gamma})$ around $\gamma = 0$:

$$M(1-\sqrt{\gamma})$$

$$= \sum_{x,y} \mathbb{P}_g(x,y|p) \exp\left(-\sqrt{\gamma} \ln\left(\frac{\mathbb{P}_g(x,y|p)}{\mathbb{P}_i(x,y|p)}\right)\right) \quad (29)$$

$$\overset{(a)}{=} \sum_{x,y} \mathbb{P}_g(x,y|p) \left(1 - \sqrt{\gamma} \ln\left(\frac{\mathbb{P}_g(x,y|p)}{\mathbb{P}_i(x,y|p)}\right) + O(\gamma)\right) \quad (30)$$

$$= 1 - \sqrt{\gamma} \sum_{x,y} \mathbb{P}_g(x,y|p) \ln\left(\frac{\mathbb{P}_g(x,y|p)}{\mathbb{P}_i(x,y|p)}\right) + O(\gamma) \quad (31)$$

$$= 1 - \sqrt{\gamma} I(X_1; Y|P=p) \ln 2 + O(\gamma). \quad (32)$$

Here (a) follows from the fact that if $\mathbb{P}_g(x,y|p) = 0$, the factor $\mathbb{P}_g(x,y|p)$ in front of the exponentiation would already cause this term to be 0, while if $\mathbb{P}_g(x,y|p) > 0$, then also $\mathbb{P}_i(x,y|p) > 0$ and thus the ratio is bounded and does not depend on γ. Substituting the above result in the original equation for ℓ we thus get the result of (28):

$$\ell = \frac{\sqrt{\gamma}(1+\sqrt{\gamma})}{-\ln M(1-\sqrt{\gamma})} \ln\left(\frac{n}{\varepsilon_1}\right) \quad (33)$$

$$= \frac{\sqrt{\gamma}(1+\sqrt{\gamma})}{\sqrt{\gamma} I(X_1; Y|P=p) \ln 2 + O(\gamma)} \ln\left(\frac{n}{\varepsilon_1}\right) \quad (34)$$

$$= \frac{\log_2 n}{I(X_1; Y|P=p)}[1 + O(\sqrt{\gamma})]. \quad (35)$$

Since the capacities tell us that $\ell / \log_2(n) \gtrsim I(X_1; Y|P=p)^{-1}$, it follows that ℓ asymptotically achieves capacity. □

Since this construction is asymptotically optimal regardless of p, in the fully informed setting we can now simply optimize p (using Theorem 1) to get the following results.

Corollary 3. *Using the values for p of Theorem 1, the asymptotics for ℓ for the five attacks of Section 2.3 are:*

$$\ell(\boldsymbol{\theta}_{int}) = 2c^2 \ln(n)\,[1 + O(\sqrt{\gamma})]\,, \quad (36)$$

$$\ell(\boldsymbol{\theta}_{all1}) = \frac{c}{\ln(2)^2} \ln(n)\,[1 + O(\sqrt{\gamma})]\,, \quad (37)$$

$$\ell(\boldsymbol{\theta}_{maj}) = \pi c \ln(n)\,[1 + O(\sqrt{\gamma})]\,, \quad (38)$$

$$\ell(\boldsymbol{\theta}_{min}) = \frac{c}{\ln(2)^2} \ln(n)\,[1 + O(\sqrt{\gamma})]\,, \quad (39)$$

$$\ell(\boldsymbol{\theta}_{coin}) = \frac{4c}{\ln(2)^2} \ln(n)\,[1 + O(\sqrt{\gamma})]\,. \quad (40)$$

Since the all-1 attack is equivalent to group testing, we mention this result separately, together with a more explicit expression for g.

Corollary 4. *Let $\boldsymbol{\theta} = \boldsymbol{\theta}_{all1}$ and let $p \approx \ln(2)/c$ be fixed. Then the log-likelihood score function g is given by*[2]

$$g(x,y) = \begin{cases} +1 & (x,y) = (0,0) \\ -1 + O(1/c) & (x,y) = (0,1) \\ -\infty & (x,y) = (1,0) \\ +c & (x,y) = (1,1) \end{cases} \quad (41)$$

Using this score function in combination with the parameters η and ℓ of Theorem 3, we obtain a simple group testing algorithm with an asymptotic number of group tests of

$$\ell \sim \frac{c \ln n}{\ln(2)^2} \approx 2.08 c \ln n, \quad (42)$$

thus achieving the simple group testing capacity.

4.1.2 Partially informed

Since the score functions from Section 4.1.1 achieve capacity for each value of p, using this score function we also trivially achieve the partially informed capacity when the arcsine distribution is used. Estimates of these capacities, and thus the resulting code lengths, can be found in Section 3.1.2.

4.1.3 Uninformed

We now arrive at what is arguably one of the most important results of this paper. Just like Oosterwijk et al. [29], who specifically studied the score function h tailored against the interleaving attack, we now also take a closer look at the log-likelihood score function designed against the interleaving attack. [3] Working out the details,

[2] To be precise: $g(0,1) = c\log_2(2 - 2^{1/c})$, and for convenience we have scaled g by a factor $c \ln 2$.

[3] Considering the interleaving attack for designing a universal decoder is further motivated by the results of Abbe and Zheng [1, 24], who showed that under certain conditions, the worst-case attack decoder is a universal capacity-achieving decoder. The interleaving attack is theoretically not the worst-case attack for finite c, but since it is known to be the asymptotic worst-case attack, the difference between the worst-case attack and the interleaving attack vanishes for large c.

this score function is of the form:

$$g(x,y,p) = \begin{cases} \ln\left(1+\frac{p}{c(1-p)}\right) & x=y=0 \\ \ln\left(1-\frac{1}{c}\right) & x \neq y \\ \ln\left(1+\frac{1-p}{cp}\right) & x=y=1 \end{cases} \tag{43}$$

The first thing to note here is that if we denote Oosterwijk et al.'s [30] score function by h, then g satisfies

$$g(x,y,p) = \ln\left(1+\frac{h(x,y,p)}{c}\right). \tag{44}$$

If $h(x,y,p) = o(c)$, then by Tayloring the logarithm around $c=\infty$ we see that $g \approx h/c$. Since scaling a score function by a constant does not affect its performance, this implies that g and h are then equivalent. Since for Oosterwijk et al.'s score function one generally needs to use *cut-offs* on F that guarantee that $h(x,y,p) = o(c)$ (cf. [16]), and since the decoder of Oosterwijk et al. is known to asymptotically achieve the uninformed capacity, we immediately get the following result.

Proposition 3. *The score function g of (43) asymptotically achieves the uninformed simple capacity when the same cut-offs on F as those in [16] are used.*

So optimizing the decoder so that it is resistant against the interleaving attack again leads to a decoder that is resistant against arbitrary attacks.

Cutting off the cut-offs Although Proposition 3 is already a nice result, we can do even better. We can prove a stronger statement, which shows one of the reasons why the log-likelihood decoder is probably more practical than the decoder of Oosterwijk et al.

Theorem 4. *The score function g of (43) achieves the uninformed simple capacity when* ***no cut-offs*** *are used.*

sketch. First note that in the limit of large c, the cut-offs of Ibrahimi et al. converge to 0. So for large c, the difference between not using cut-offs and using cut-offs is negligible, as long as the contribution of the tails of p near 0 or 1 to the distribution of user scores is negligible. Since with this score function g, all moments of both innocent and guilty user scores are finite (arbitrary powers of logarithms always lose against the $1/\sqrt{p(1-p)}$ of the arcsine distribution and the decreasing width of the interval between 0 and the cut-off), the tails indeed decay exponentially. So also without cut-offs this score function asymptotically achieves the uninformed simple capacity. □

Note that the same result does not apply to the score function of Oosterwijk et al. [30], for which the tails of the distributions are not Gaussian enough to omit the use of cut-offs. The main difference is that for small p, the score function h of [29] scales as $1/p$ (which explodes when p is really small), while the log-likelihood decoder g then only scales as $\ln(1/p)$ which is much smaller.

All roads lead to Rome Let us now mention a third way to obtain a capacity-achieving uninformed simple decoder which is again very similar to the two decoders above. To construct this decoder, we use a Bayesian approximation of the proposed empirical mutual information decoder of Moulin [25], and again plug in the asymptotic worst-case attack, the interleaving attack.

Theorem 5. *Using Bayesian inference with an a priori probability of guilt of $\mathbb{P}(j \in \mathcal{C}) = \frac{c}{n}$, the empirical mutual information decoder tailored against the interleaving attack can be approximated with the following score function:*

$$m(x,y,p) = \begin{cases} \ln\left(1+\frac{p}{n(1-p)}\right) & x=y=0 \\ \ln\left(1-\frac{1}{n}\right) & x \neq y \\ \ln\left(1+\frac{1-p}{np}\right) & x=y=1 \end{cases} \tag{45}$$

sketch. For now, let p be fixed. The empirical mutual information decoder assigns a score S_j to a user j using

$$S_j = \sum_{x,y} \hat{\mathbb{P}}(x,y) \ln\left(\frac{\hat{\mathbb{P}}(x,y)}{\hat{\mathbb{P}}(x)\hat{\mathbb{P}}(y)}\right) = \sum_{i=1}^{\ell} \ln\left(\frac{\hat{\mathbb{P}}(x_{j,i},y_i)}{\hat{\mathbb{P}}(x_{j,i})\hat{\mathbb{P}}(y_i)}\right) \tag{46}$$

where $\hat{\mathbb{P}}(\cdot)$ denotes the empirical estimate of $\mathbb{P}(\cdot)$ based on the data $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{p}$. For large ℓ, these estimates will converge to the real probabilities, so we can approximate S_j by

$$S_j \approx \sum_{i=1}^{\ell} \ln\left(\frac{\mathbb{P}(x_{j,i},y_i)}{\mathbb{P}(x_{j,i})\mathbb{P}(y_i)}\right) = \sum_{i=1}^{\ell} m(x_{j,i},y_i,p_i). \tag{47}$$

Here $\mathbb{P}(X_{j,i})$ and $\mathbb{P}(y_i)$ can be easily computed, but for computing $\mathbb{P}(X_{j,i},y_i)$ we need to know whether user j is guilty or not. Using Bayesian inference, we can write

$$\mathbb{P}(x,y) = \mathbb{P}_g(x,y)\mathbb{P}(j \in \mathcal{C}) + \mathbb{P}_i(x,y)\mathbb{P}(j \notin \mathcal{C}). \tag{48}$$

Assuming an a priori probability of guilt of $\mathbb{P}(j \in \mathcal{C}) = c/n$, we can work out the details to obtain

$$m(x,y,p) = \ln\left(1+\frac{c}{n}\left[\frac{\mathbb{P}_g(x,y)}{\mathbb{P}_i(x,y)}-1\right]\right). \tag{49}$$

Filling in the corresponding probabilities for the interleaving attack, we end up with the score function of (45). □

For values of p with $o(1) < p < 1-o(1)$, this decoder is again equivalent to both the log-likelihood score function g and Oosterwijk et al.'s score function h.

4.2 Joint decoders

For the joint decoding setting, scores are assigned to tuples of c users, and again higher scores correspond to a higher probability of being accused. The most natural step from the simple log-likelihood decoders to joint decoders seems to be to use the following joint score function:

$$g(x_1,\ldots,x_c,y,p) = \ln\left(\frac{\mathbb{P}_{g^c}(x_1,\ldots,x_c,y|p)}{\mathbb{P}_{i^c}(x_1,\ldots,x_c,y|p)}\right). \tag{50}$$

Here $\mathbb{P}_{g^c}(\cdot)$ is under the assumption that in this tuple *all users are guilty*, while for $\mathbb{P}_{i^c}(\cdot)$ we assume that *all users are innocent*. Note that under the assumption that the attack is colluder-symmetric, the score function only depends on $z = \sum_{i=1}^{c} x_i$:

$$g(x_1, \ldots, x_c, y, p) = g(z, y, p) = \ln\left(\frac{\mathbb{P}_{g^c}(z, y|p)}{\mathbb{P}_{i^c}(z, y|p)}\right). \quad (51)$$

4.2.1 Fully informed

To analyze the joint decoder, we again make use of the moment generating function for the score assigned to tuples of c innocent users. This function is now defined by

$$M(t) = \sum_{z,y} \mathbb{P}_{i^c}(z, y|p)^{1-t} \mathbb{P}_{g^c}(z, y|p)^t \quad (52)$$

and it satisfies $M(t) = \mathbb{E}_{i^c} e^{tS_{j,i}} = \mathbb{E}_{g^c} e^{(t-1)S_{j,i}}$ and $M(0) = M(1) = 1$. Using similar techniques as in Section 4.1.1, we obtain the following result.

Theorem 6. *Let p and $\boldsymbol{\theta}$ be fixed and known to the distributor. Let $\gamma = \ln(1/\varepsilon_2)/\ln(n^c/\varepsilon_1)$, and let the threshold η and code length ℓ be defined as*

$$\eta = \ln\left(\frac{n^c}{\varepsilon_1}\right), \qquad \ell = \frac{\sqrt{\gamma}(1+\sqrt{\gamma})}{-\ln M(1-\sqrt{\gamma})} \ln\left(\frac{n^c}{\varepsilon_1}\right). \quad (53)$$

Then with probability at least $1-\varepsilon_1$ all all-innocent tuples are not accused, and with probability at least $1-\varepsilon_2$ the single all-guilty tuple is accused.

sketch. The proof is very similar to the proof of Theorem 3. Instead of n innocent and c guilty users we now have $\binom{n}{c} < n^c$ all-innocent tuples and just 1 all-guilty tuple, which changes some of the numbers in γ, η and ℓ. We again apply the Markov inequality with $\alpha = 1$ for innocent tuples and $\beta = \sqrt{\gamma}$ for guilty tuples, to obtain the given expressions for η and ℓ. □

Note that Theorem 6 does not prove that we can actually find the set of colluders with high probability, since mixed tuples consisting of both innocent and guilty users also exist, and these may or may not have a score exceeding η. This does prove that with high probability we can find a set $\mathcal{C}'$ of c users, for which (i) all tuples not containing these users have a score below η, and (ii) the tuple containing exactly these users has a score above η. Regardless of what the scores for mixed tuples are, with probability at least $1-\varepsilon_1-\varepsilon_2$ such a set consists and contains at least one colluder. Furthermore, if this set $\mathcal{C}'$ is unique, then with high probability this is exactly the set of colluders. But there is no guarantee that it is unique without additional proofs. This is left for future work.

To further motivate why using this joint decoder may be the right choice, the following proposition shows that at least the scaling of the resulting code lengths is optimal. Note that the extra c that we get from $\ln(n^c) = c \ln n$ can be combined with the mutual information $I(\cdot)$ to obtain $\frac{1}{c}I(\cdot)$, which corresponds to the joint capacity.

Proposition 4. *If $\gamma = o(1)$ then the code length ℓ of Theorem 6 scales as*

$$\ell = \frac{\log_2 n}{\frac{1}{c}I(Z; Y|P=p)} \left[1 + O(\sqrt{\gamma})\right], \quad (54)$$

thus asymptotically achieving the optimal code length (up to first order terms) for arbitrary values of p.

Since the asymptotic code length is optimal regardless of p, these asymptotics are also optimal when p is optimized to maximize the mutual information in the fully informed setting.

Finally, although it is hard to estimate the scores of mixed tuples with this decoder, just like in [31] we expect that the joint decoder score for a tuple is roughly equal to the sum of the c individual simple decoder scores. So a tuple of c users consisting of k colluders and $c-k$ innocent users is expected to have a score roughly a factor k/c smaller than the expected score for the all-guilty tuple. So after computing the scores for all tuples of size c, we can get rough estimates of how many guilty users are contained in each tuple, and for instance try to find the set $\mathcal{C}'$ of c users that best matches these estimates. There are several options for post-processing that may improve the accuracy of using this joint decoder, which are left for future work.

4.2.2 Partially informed

As mentioned in Proposition 4, the code length is asymptotically optimal regardless of p, so the code length in the partially uninformed setting is also asymptotically optimal. Asymptotics on ℓ can thus be obtained by combining Proposition 4 with the results of Section 3.2.2.

4.2.3 Uninformed

Note that if the above joint decoder turns out to work well, then we can again plug in the interleaving attack to get something that might just work well against arbitrary attacks. While we cannot prove that this joint decoder is optimal, we can already see what the score function would be, and conjecture that it works against arbitrary attacks.

Conjecture 1. *The joint log-likelihood decoder against the interleaving attack, with the score function g defined by*

$$g(z, y, p) = \begin{cases} \ln(1 - \frac{z}{c}) - \ln(1-p) & (y=0) \\ \ln(\frac{z}{c}) - \ln(p) & (y=1) \end{cases} \quad (55)$$

works against arbitrary attacks and asymptotically achieves the joint capacity of the uninformed fingerprinting game.

A further study of this universal joint decoder is left as an open problem.

5 Discussion

Let us now briefly discuss the results from Sections 3 and 4, their consequences, and some directions for future work.

Informed simple decoding For the setting of simple decoders, we derived explicit asymptotics on the informed capacities for various attacks, which often scale as $\Theta(c^{-1})$. We further showed that log-likelihood scores provably match these bounds for large n, regardless of $\boldsymbol{\theta}$ and p. Because these decoders are optimal for any value of p, they are also optimal in the partially informed setting, where different values of p are used. If the encoder uses the arcsine distribution to generate biases, we showed that these capacities generally seem to scale as $\Theta(c^{-3/2})$, which is roughly 'halfway' between the fully informed and uninformed capacities.

Uninformed simple decoding Although log-likelihood decoders have already been studied before in the context of fingerprinting, the main drawback was always that to use these decoders, you would either have to fill in (and know) the exact pirate strategy, or compute the worst-case attack explicitly. So if you are in the simple uninformed setting where you don't know the pirate strategy and where the worst-case attack is not given by a nice closed-form expression [14, Fig. 4b], how can you construct such decoders for large c? The trick seems to be to just fill in the *asymptotic* worst-case attack, which Huang and Moulin showed is the interleaving attack [14], and which is much simpler to analyze. After previously suggesting this idea to Oosterwijk et al., we now used the same trick here to obtain two other capacity-achieving score functions using two different methods (but each time filling in the interleaving attack). So in total we now have three different methods to obtain (closed-form) capacity-achieving decoders in the uninformed setting:

- Using Lagrange-multipliers, Oosterwijk et al. [29] obtained:

$$h(x,y,p) = \begin{cases} +\frac{p}{1-p} & x=y=0 \\ -1 & x \neq y \\ +\frac{1-p}{p} & x=y=1 \end{cases} \tag{56}$$

- Using Neyman-Pearson-based log-likelihood scores, we obtained:

$$g(x,y,p) = \begin{cases} \ln\left(1+\frac{p}{c(1-p)}\right) & x=y=0 \\ \ln\left(1-\frac{1}{c}\right) & x \neq y \\ \ln\left(1+\frac{1-p}{cp}\right) & x=y=1 \end{cases} \tag{57}$$

- Using a Bayesian approximation of the empirical mutual information decoder of Moulin [25], we obtained:

$$m(x,y,p) = \begin{cases} \ln\left(1+\frac{p}{n(1-p)}\right) & x=y=0 \\ \ln\left(1-\frac{1}{n}\right) & x \neq y \\ \ln\left(1+\frac{1-p}{np}\right) & x=y=1 \end{cases} \tag{58}$$

For $o(1) < p < 1-o(1)$ and large c, n, these score functions are equivalent up to a scaling factor:

$$h(x,y,p) \sim c \cdot g(x,y,p) \sim n \cdot m(x,y,p), \tag{59}$$

and therefore all three are asymptotically optimal. So there may be many different roads that lead to Rome, but they all seem to have one thing in common: to build a universal decoder that works against arbitrary attacks, one should build a decoder that works against the asymptotic worst-case pirate attack, the interleaving attack. And if it does work against this attack, then it probably works against any other attack as well.

Joint decoding Although deriving the joint informed capacities is much easier than deriving the simple informed capacities, actually building decoders that provably match these bounds is a different matter. We conjectured that the same log-likelihood scores achieve capacity when a suitable accusation algorithm is used, and we conjectured that the log-likelihood score built against the interleaving attack achieves the uninformed joint capacity, but we cannot prove any of these statements beyond reasonable doubt. For now this is left as an open problem.

Group testing Since the all-1 attack is equivalent to group testing, some of the results we obtained also apply to group testing. The joint capacity was already known [33], but to the best of our knowledge both the simple capacity (Corollary 1) and a simple decoder matching this simple capacity (Corollary 4) were not yet known before. Attempts have been made to build efficient simple decoders with a code length not much longer than the joint capacity [6], but these do not match the simple capacity. Future work will include computing the capacities and building decoders for various noisy group testing models, where the marking assumption may not apply.

Dynamic fingerprinting Although this paper focused on applications to the 'static' fingerprinting game, the construction of [20] can trivially be applied to the decoders in this paper as well to build efficient dynamic fingerprinting schemes. Although the asymptotics for the code length in this dynamic construction are the same, (i) the order terms are significantly smaller in the dynamic game, (ii) one does not need the assumption that the pirate strategy is colluder-symmetric, and (iii) one does not necessarily need to know (a good estimate of) c in advance [20, Section V]. An important open problem remains to determine the dynamic uninformed fingerprinting capacity, which may prove or disprove that the construction of [20] is optimal.

Further generalizations While this paper already aims to provide a rather complete set of guidelines on what to do in the various different fingerprinting games (with different amounts of side-information, and different computational assumptions on the decoder), there are some further generalizations that were not considered here due to lack of space. We mention two in particular:

- **Larger alphabets**: In this work we focused on the binary case of $q = 2$ different symbols, but it may be advantageous to work with larger alphabet sizes $q > 2$, since the code length decreases linearly with q. For

the results about decoders we did not really use that we were working with a binary alphabet, so it seems a straightforward exercise to prove that the q-ary versions of the log-likelihood decoders also achieve capacity. A harder problem seems to be to actually compute these capacities in the various informed settings, since the maximization problem then transforms from a one-dimensional optimization problem to a $(q-1)$-dimensional optimization problem.

- **Tuple decoding**: As in [31], we can consider a setting in between the simple and joint decoding settings, where decisions to accuse are made based on looking at tuples of users of size at most t. Tuple decoding may offer a trade-off between the high complexity, low code length of a joint decoder and the low complexity, higher code length of a simple decoder, and so it may be useful to know how the capacities scale in the region $1 < t < c$.

6 Acknowledgments

The author is very grateful to Pierre Moulin for his insightful comments and suggestions during the author's visit to Urbana-Champaign that inspired work on this paper. The author would also like to thank Teddy Furon for pointing out the connection between decoders designed against the interleaving attack and the results of Abbe and Zheng [1], and for finding some mistakes in a preliminary version of this manuscript. Finally, the author thanks Jeroen Doumen, Jan-Jaap Oosterwijk, Boris Škorić, and Benne de Weger for valuable discussions and comments.

References

[1] E. Abbe and L. Zheng. Linear universal decoding for compound channels. *IEEE Transactions on Information Theory*, 56(12):5999–6013, 2010.

[2] E. Amiri and G. Tardos. High rate fingerprinting codes and the fingerprinting capacity. In *20th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 336–345, 2009.

[3] O. Blayer and T. Tassa. Improved versions of Tardos' fingerprinting scheme. *Designs, Codes and Cryptography*, 48(1):79–103, 2008.

[4] D. Boesten and B. Škorić. Asymptotic fingerprinting capacity for non-binary alphabets. In *13th Conference on Information Hiding (IH)*, pages 1–13, 2011.

[5] D. Boneh and J. Shaw. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory*, 44(5):1897–1905, 1998.

[6] C.-L. Chan, S. Jaggi, V. Saligrama, and S. Agnihotri. Non-adaptive group testing: explicit bounds and novel algorithms. In *IEEE International Symposium on Information Theory (ISIT)*, pages 1837–1841, 2012.

[7] A. Charpentier, F. Xie, C. Fontaine, and T. Furon. Expectation maximization decoding of Tardos probabilistic fingerprinting code. In *SPIE Proceedings*, volume 7254, 2009.

[8] T. M. Cover and J. A. Thomas. *Elements of Information Theory (2nd Edition)*. Wiley Press, 2006.

[9] A. Fiat and T. Tassa. Dynamic traitor tracing. *Journal of Cryptology*, 14(3):354–371, 2001.

[10] T. Furon and L. Pérez-Freire, EM decoding of Tardos traitor tracing codes. In *ACM Symposium on Multimedia and Security (MM&Sec)*, pages 99–106, 2009.

[11] Y.-W. Huang and P. Moulin. Capacity-achieving fingerprint decoding. In *IEEE Workshop on Information Forensics and Security (WIFS)*, pages 51–55, 2009.

[12] Y.-W. Huang and P. Moulin. Saddle-point solution of the fingerprinting capacity game under the marking assumption. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2256–2260, 2009.

[13] Y.-W. Huang and P. Moulin. Maximin optimality of the arcsine fingerprinting distribution and the interleaving attack for large coalitions. In *IEEE Workshop on Information Forensics and Security (WIFS)*, pages 1–6, 2010.

[14] Y.-W. Huang and P. Moulin. On the saddle-point solution and the large-coalition asymptotics of fingerprinting games. *IEEE Transactions on Information Forensics and Security*, 7(1):160–175, 2012.

[15] Y.-W. Huang and P. Moulin. On fingerprinting capacity games for arbitrary alphabets and their asymptotics. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2571–2575, 2012.

[16] S. Ibrahimi, B. Škorić, and J.-J. Oosterwijk. Riding the saddle point: asymptotics of the capacity-achieving simple decoder for bias-based traitor tracing. *Cryptology ePrint Archive*, 2013.

[17] T. Laarhoven, J.-J. Oosterwijk, and J. Doumen. Dynamic traitor tracing for arbitrary alphabets: divide and conquer. In *IEEE Workshop on Information Forensics and Security (WIFS)*, pages 240–245, 2012.

[18] T. Laarhoven and B. de Weger. Discrete distributions in the Tardos scheme, revisited. In *1st ACM Workshop on Information Hiding and Multimedia Security (IH&MMSec)*, pages 13–18, 2013.

[19] T. Laarhoven. Efficient probabilistic group testing based on traitor tracing. In *51st Annual Allerton Conference on Communication, Control and Computing (Allerton)*, 2013.

[20] T. Laarhoven, J. Doumen, P. Roelse, B. Škorić, and B. de Weger. Dynamic Tardos traitor tracing schemes. *IEEE Transactions on Information Theory*, 59(7):4230–4242, 2013.

[21] T. Laarhoven. Dynamic traitor tracing schemes, revisited. In *IEEE Workshop on Information Forensics and Security (WIFS)*, pages 191–196, 2013.

[22] T. Laarhoven and B. de Weger. Optimal symmetric Tardos traitor tracing schemes. *Designs, Codes and Cryptography*, 71(1): 83–103, 2014.

[23] P. Meerwald and T. Furon. Group testing meets traitor tracing. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4204–4207, 2011.

[24] P. Meerwald and T. Furon. Toward practical joint decoding of binary Tardos fingerprinting codes. *IEEE Transactions on Information Forensics and Security*, 7(4):1168–1180, 2012.

[25] P. Moulin. Universal fingerprinting: capacity and random-coding exponents. *arXiv:0801.3837v3 [cs.IT]*, 2011.

[26] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 231:289–337, 1933.

[27] K. Nuida, M. Hagiwara, H. Watanabe, and H. Imai. Optimization of Tardos's fingerprinting codes in a viewpoint of memory amount. In *9th Conference on Information Hiding (IH)*, pages 279–293, 2007.

[28] K. Nuida, S. Fujitsu, M. Hagiwara, T. Kitagawa, H. Watanabe, K. Ogawa, and H. Imai. An improvement of discrete Tardos fingerprinting codes. *Designs, Codes and Cryptography*, 52(3):339–362, 2009.

[29] J.-J. Oosterwijk, B. Škorić, and J. Doumen. Optimal suspicion functions for Tardos traitor tracing schemes. In *1st ACM Workshop on Information Hiding and Multimedia Forensics (IH&MMSec)*, pages 19–28, 2013.

[30] J.-J. Oosterwijk, B. Škorić, and J. Doumen. A capacity-achieving simple decoder for bias-based traitor tracing schemes. *Cryptology ePrint Archive*, 2013.

[31] J.-J. Oosterwijk, J. Doumen, and T. Laarhoven. Tuple decoders for traitor tracing schemes. In *SPIE Proceedings*, volume 9028, 2014.

[32] L. Pérez-Freire and T. Furon. Blind decoder for binary probabilistic traitor tracing codes. In *IEEE Workshop on Information Forensics and Security (WIFS)*, pages 46–50, 2009.

[33] A. Sebő. On two random search problems. *Journal of Statistical Planning and Inference*, 11:23–31, 1985.

[34] A. Simone and B. Škorić. Accusation probabilities in Tardos codes: beyond the Gaussian approximation. *Designs, Codes and Cryptography*, 63(3):379–412, 2012.

[35] B. Škorić, T. U. Vladimirova, M. U. Celik, and J. C. Talstra. Tardos fingerprinting is better than we thought. *IEEE Transactions on Information Theory*, 54(8):3663–3676, 2008.

[36] B. Škorić, S. Katzenbeisser, and M. U. Celik. Symmetric Tardos fingerprinting codes for arbitrary alphabet sizes. *Designs, Codes and Cryptography*, 46(2):137–166, 2008.

[37] B. Škorić and J.-J. Oosterwijk. Binary and q-ary Tardos codes, revisited. *Designs, Codes and Cryptography*, 2013.

[38] G. Tardos. Optimal probabilistic fingerprint codes. In *35th ACM Symposium on Theory of Computing (STOC)*, pages 116–125, 2003.